# PERIODIC INTERFACE CALIBRATION FOR HIGH SPEED COMMUNICATION

Inventor: Jade Kizer

## BACKGROUND OF THE INVENTION

Field of the Invention

**[0001]** The present application relates to high-speed communication interfaces, including high-speed parallel bus interfaces for integrated circuits; and more particularly to calibration of such interfaces.

Description of Related Art

**[0002]** High-performance data processing applications are driving the demand for data rates past the GigaHertz range. As processor clock speeds increase to meet the demand, high-performance parallel bus interface technology is being developed to meet these needs. In parallel bus interfaces, a number of serial lines are operated in parallel. So-called SERDES (short for serializer-deserializer) technologies are being applied for each of the parallel lines. Other high-performance bus interface technologies are provided by Rambus, Inc., including products provided under the tradenames XDR™ High Performance Memory Interface Technology, Raser™ High Performance Interface Technology, and Redwood™ High Performance Parallel Bus Interface Technology. Background concerning high speed interfaces is found in U.S. Patent No. 6,396,329 B1, entitled Method and Apparatus for Receiving High Speed Signals with Low Latency; and in U.S. Patent No.6,473,439, entitled Method and Apparatus for Fail-Safe Resynchronization with Minimum Latency.

**[0003]** One problem which becomes more important as communication speeds increase is calibration of clocks and sample timing. The optimal sampling point for each bit of data is controlled by many independent variables, which can be boiled down to a simple relationship between clock and data. There is an optimal singular sampling point for all data patterns at any given moment. Complicating matters are changes to the optimal sampling point. High-frequency noise, known as jitter, places a cloud of uncertainty around this optimal sampling point. Methods to compensate for jitter have

been limited in effectiveness. Thus, systems with very low jitter are preferred. Low-frequency noise, known as skew, comprises slowly changing offsets in the optimal sampling point, for which compensation can be provided, depending on the system's ability to track of these sources of error.

**[0004]** Several methods have been developed to track and calibrate the sources of error that cause skew. One method is known as oversampling. Oversampling requires sampling the data more than once per bit time and coding the data for guaranteed transitions. These oversampling approaches involve clock/data recovery schemes that use clock/data patterns such as 8b/10b, and the like. Most current SERDES technologies use the 8b/10b coding scheme. This approach has the advantage that it relies on the same number of physical channels as logical channels for the communication link. However, there is an inherent 25% bandwidth penalty built-in the 8b/10b coding scheme. Also, the oversampling requires increased power consumption.

**[0005]** Another method for tracking and calibrating sources of error of involves performing an initial calibration, and then letting the system run open loop. This process requires good circuits to track all temperature-related drift components. One well-known example of this approach is known as the source synchronous technique. A timing reference is sent, typically on an independent physical channel, along with the data to compensate for drift between clock and data. The tracking time constant needs to be as fast as possible, with minimum time lag. Additionally, a single offset value would be optimal for all operating conditions on each of the lines in the parallel bus. If good tracking can be achieved across all drift conditions on all of the lines in the parallel bus, a source synchronous approach is quite compelling.

**[0006]** In another approach, where tracking times are not optimal, each link can be temporarily disabled and used for a fast periodic calibration. This type of periodic calibration requires precise logical synchronization between transmit and receive operations to perform the calibration efficiently during a calibration window, without jeopardizing real data in the process. Although synchronized periodic operations may be possible in a master-slave implementation, peer-to-peer periodic operations may be too prohibitive to be efficiently incorporated.

**[0007]** Selecting an optimal chip-to-chip interconnect strategy relies not only on the traditional metrics of latency and effective bandwidth, but also the area and power required to do so. System solutions that provide superior area/bandwidth and power/bandwidth trade-offs, while still meeting the bandwidth and latency requirements of system designers, are required to continue to scale performance in line with expected trends.

## SUMMARY OF THE INVENTION

**[0008]** The present invention provides a communications interface, including transmitters and receivers, adapted for periodic calibration, and a method for maintaining calibration of communication paths across the interface. The periodic calibration process can operate substantially continuously, as a background process during operation of the interface to maintain the communication lines during long intervals of constant use without reset or other initialization events that allow time for typical line maintenance operations.

**[0009]** A method according to the present invention manages a high-speed communication interface for a parallel bus having N bus lines at the logical layer. In the physical layer, N+1 communication lines are established. A maintenance operation (calibration for example) is performed on one of the N+1 communication lines, while N of the N+1 communication lines is available for data from the N line parallel bus. The communication line on which the maintenance operation is performed, is changed after the operation is complete, so that all of the N+1 communication lines are periodically maintained, without interfering with communications on N of the N+1 communication lines.

**[0010]** Where the maintenance operation is calibration, a calibration signal, such as a pseudorandom bit sequence adapted for calibration of receiver clocks, is transmitted from a source, and received at a destination, on one particular communication line, referred to as communication line (n), of the N+1 communication lines. At the same time, a path is maintained for communication of data on N communication lines. A parameter associated with communication line (n) is calibrated. Then, after calibrating the parameter associated with communication line (n), the index (n) is changed and the

process is repeated for a next communication line. Accordingly, one of the N+1 communication lines is used for calibration at a time, and is rotated according to a pattern so that each of the N+1 communication lines is calibrated over time.

**[0011]** Embodiments of the method include entering a reduced power consumption state on at least one of receivers and transmitters on the N communication lines, for example when data is not being supplied from the N line bus, while continuing to perform the periodic maintenance operation on the communication lines. In this manner, powerdown states are supported without losing maintenance, such as calibration, of the high-speed parallel data interface. In one embodiment, the system includes a "nap" mode, during which the periodic maintenance procedure continues, while other circuitry supporting the communication lines are in a power down state. During the "nap" mode, the maintenance procedure may operate with a cycle time that is less than, the same as or greater than the cycle time during normal operations of the communication lines. The system may also support a mode in which the maintenance procedures are stopped, and circuitry supporting the maintenance procedures is in a power down state.

**[0012]** The present invention also provides a method for switching the communication line subject of maintenance, without interrupting dataflow. The method includes, for example, changing the index (n) to switch a first particular communication line from being subject of maintenance to communicating from a line on the N line bus, and a second particular communication line from communicating from the line on the N line bus to being subject of maintenance, routing the first and second particular communication lines together from the line in the N line bus during a settling interval, and then, after the settling interval, performing maintenance on the second particular communication line.

**[0013]** The index (n) is changed in embodiments of the present invention according to a continuous periodic function, so that each of the N+1 communication lines is maintained at least once during a period of the continuous periodic function. Where the set of N+1 communication lines includes communication lines logically identified as paths 0 to N, one pattern comprises a repeating pattern beginning with the index (n) equal to zero, and increasing to (n) equal to N, and then decreasing to (n) equal to zero.

**[0014]** The present invention is also embodied by signal interfaces supporting the source and destination ends of the communication lines. Thus, an embodiment of the invention includes a set of signal lines having N+1 signal lines and N+1 receivers coupled to respective signal lines in a set of signal lines, which together establish a set of N+1 signal paths. The set of N+1 signal paths is adapted to serve an N line bus. A line maintenance circuit, such as the calibration circuit, is included in the interface. A switch placed in the N+1 signal paths, such as between the N+1 receivers and the N line bus, and control logic for the switch, operate to selectively route N signal paths in the set to the N line bus and one signal path, signal path (n), in the set to the line maintenance circuit. The index (n) is changed as discussed above to maintain the signal paths in the set without interfering with dataflow.

**[0015]** According to embodiments of the invention, the line maintenance circuit comprises a calibration circuit. For example, a calibration circuit is used to set adjustable clock generators that are used to supply receiver clocks for each of the N+1 receivers.

**[0016]** In yet other embodiments, logic is included to power down the N+1 receivers while continuing to maintain signal paths in a set of signal paths according to the pattern.

**[0017]** Other embodiments of the invention are implemented on the source side of the communication line. In such embodiments, an N line bus feeds a set of signal lines having N+1 signal lines. N+1 transmitters are coupled to the set of signal lines establishing a set of N+1 signal paths. A line maintenance circuit is included. A switch is coupled to the N+1 signal paths. Control logic for the switch selectively routes N signal paths in the set from the N line bus to N signal lines in the set of signal lines, and routes one signal path, signal path (n), in the set from the line maintenance circuit to the signal line (n). The index (n) is changed so that the line maintenance circuit is successively coupled to each of the N+1 signal lines according to the pattern, such as described above. Likewise, on the source side of the communication line, the N+1 transmitters can be powered down without interfering with the line maintenance process.

**[0018]** Further embodiments of the invention comprise the combination of the source side, destination side and communication media to provide a complete high-speed, parallel communication system.

**[0019]** Embodiments of the present invention support communication between integrated circuits at data rates greater than 100MHz, and in some embodiments greater than 1GHz, and more.

**[0020]** The example of line maintenance mentioned above involves transmission of calibration signals used for example for calibration of receiver clocks. Calibration can be applied to other parameters of the communication line, such as signaling levels, optimal placement of sampling times for symbol capture, and impedance of the termination element, and receiver thresholds. For communication lines that use adaptive equalization or filters, the maintenance can include adjustment of equalization or filter coefficients. The line management process can be applied to line maintenance applications which may or may not involve transmission of calibration signals.

**[0021]** Other aspects and advantages of the present invention can be seen on review of the drawings, the detailed description and the claims, which follow.

## BRIEF DESCRIPTION OF THE DRAWINGS

**[0022]** Fig. 1 is a simplified block diagram of a system employing periodic calibration.

**[0023]** Fig. 2 is a more detailed diagram of physical layer signal paths in a system using continuous periodic calibration.

**[0024]** Fig. 3 shows flowcharts for continuous periodic calibration on the transmit side and on the receive side.

**[0025]** Fig. 4 illustrates circuitry for calibration of a receive clock in a circuit such as shown in Fig. 2.

**[0026]** Fig. 5 is a simplified block diagram of a system employing periodic calibration, in combination with a source synchronous clock.

## DETAILED DESCRIPTION

**[0027]** A detailed description of embodiments of the present invention is provided with reference to Figs. 1-5.

**[0028]** Fig. 1 is a simplified block diagram of the communication system applying continuous periodic calibration according to the present invention. The system includes a

first integrated circuit 10 and a second integrated circuit 11. The first integrated circuit 10 includes a logical layer parallel bus 20 including N lines, a calibration signal source 21, and calibration logic 22. A switch 23 couples the parallel bus 20 and the calibration signal source 21 with a set of transmitters 12-16, including one for each of N+1 physical layer communication lines. The set of transmitters 12-16 drives communication signals across communication media. In this example, the set of transmitters 12-16 drive data on signal lines coupled to input/output ports 32-36 (such as IO pins on the integrated circuit), which are coupled to respective transmission lines, including line 0 through line N in a set of N+1 transmission lines.

**[0029]** The second integrated circuit 11 includes complementary components. Input/output ports 42-46 are coupled by signal lines to respective receivers 52-56 in a set of N+1 receivers on the second integrated circuit 11. The receivers 52-56 are coupled to switch 57. Switch 57 routes the outputs of N receivers from the set to an N-line parallel bus 58, while routing the output of one of the receivers from the set to a calibration circuit 59. Calibration logic 60 on the second integrated circuit 11 controls the switch 57 and calibration circuit 59 to manage the continuous periodic calibration of the set of communication lines.

**[0030]** The logic 22 in the first integrated circuit 10 and the logic 60 in the second integrated circuit 11 support a nap state, in which the transmitters and receivers are placed in a power down mode when not needed, while the calibration cycle continues. This nap state maintains readiness of the high-speed parallel interface for fast transition from power conserving conditions to awake operations in the transmitting and receiving systems. In some embodiments, another low power state is included in which the calibration process is also stopped.

**[0031]** In Fig. 1, the communication links are shown operating in one direction. The invention is also extended to bidirectional communication links, where the receivers and transmitters, and other supporting logic, are found on both the first and second integrated circuits.

**[0032]** The continuous periodic calibration process uses an extra link to provide a mechanism to time multiplex the calibrating operation across an interface. In a parallel interface with eight logical links (N=8), nine physical links would be used with one

assigned in a rotating pattern to be the calibration link. The calibration link spends as much time calibrating as necessary, without affecting worst-case latency of the system. The rate of rotation among the set of communication links can be adapted to suit the needs of the particular implementation. For example, in systems applying spread spectrum clocking, in which the clock rate is varied over a relatively slow interval, the rate of rotation among the communication links should be high enough that the changes in clock rate due to spread spectrum processing are not impacted. In other examples, the rate of rotation should be fast enough to accommodate known sources of skew of the parameter being calibrated, such as temperature drift coefficients.

**[0033]** The continuous periodic calibration process can be extended to glue multiple parallel interfaces together, for example in a daisy chain configuration. The process yields a worst-case reduction in overall effective bandwidth to 1/N+1, where N is the number of logical links in the system. The rotation of the calibration link is done entirely in the physical layer in preferred embodiments, providing a seamless N link logical layer to the host system.

**[0034]** One difficulty can arise in the handoff between the rotating calibration link and a regular data link, especially in cases without any back channel communication supporting the rotation of the calibration operation. The handoff must be executed without loss of data in the transition, and requires some synchronization between the transmit side and the receive side.

**[0035]** Fig. 2 illustrates one particular implementation of the physical layer in a high-speed parallel communication interface according to the present invention. With an input bus supplying eight bits of data (N=8), eight transmit data sources TDATA[0] to TDATA[7] (100-107) provide eight inputs from the logical layer. A transmit calibration signal source TXCAL (108) provides a ninth input. A switch includes nine physical layer, three-input multiplexers 110-118 having outputs coupled to respective transmitters TX_IO[0] to TX_IO[8] (120-128), which drive data on respective communication media 130-138. The inputs to the multiplexers 110-118 can be characterized with respect to the index (n), where (n) is an integer from 0 to N corresponding to the N+1 communication media 130-138. Each multiplexer (n) has as input the output of the transmit calibration source TXCAL 108, and input bus lines TDATA[n-1] and TDATA[n], except on the

boundaries where multiplexer (0) 110 receives the input bus line TDATA[0] only, and multiplexer (N=8 to) 118 receives the input bus line TDATA[7] only. Extra inputs on the multiplexers 110 and 118 on the boundary may be used to support daisy-chaining multiple buses.

**[0036]** On the receive side, N+1 receivers RX_IO[0] to RX_IO[8] (140-148) are coupled to respective communication media 130-138. The outputs of the receivers 140-148 are each coupled to respective buffers 150-158 which drive the outputs to a receiver calibration circuit RXCAL 171. In addition, a switch includes eight physical layer, two-input multiplexers 160-167 coupled to the outputs of the receivers 140-148. The inputs to the multiplexers 160-167, can be characterized with respect to the index (n), where (n) ranges from 0 to N-1. Thus, the input two multiplexer(n) in the set of multiplexers 160-167 on the receive side include the outputs of receivers RX_IO[n] and RX_IO[n+1]. The outputs of the multiplexers 160-167 are coupled to the logical layer N line bus 170, providing a receive data RXDATA[7:0] (note that RXDATA[7:0] can be thought of as performing the reciprocal function of the TDATA[0] 100 – TDATA[7] 107).

**[0037]** Control logic associated with the multiplexers on both sides manages the handoff between changing operation of a first particular communication link for calibration to communicating data, and changing operation of a second particular communication link from communicating data to operation for calibration. In one example, the link being calibrated is rotating among the set of N+1 communication lines according to a continuous periodic pattern, as can be understood with reference to an example, as follows. Assume that the communication link corresponding to index (n=1), comprising transmitter TX_IO[1] 121, communication medium 131, and receiver RX_IO[1], is currently assigned to the calibration task. In this state, the transmit data lines TDATA[7:0] from the input bus map to transmitters TX_IO[8:2,0]. Likewise, receivers RX_IO[8:2,0] map to receive data lines RXDATA[7:0]. The transmitter TX_IO[1] is transmitting calibration data, and the receiver RX_IO[1] is coupled to the receive calibration circuit RXCAL 171 via buffer 151. After calibration is complete on the receiver RX_IO[1], such as by adjusting the clock to the calibrated sampling point, the system is ready to change the link on which calibration is executed. An operation to

switch the communication link assigned to the calibration task from index (n=1) to index (n=2) is executed as follows.

**[0038]** A) TDATA[1] is mapped to both transmitters TX_IO[1] and TX_IO[2].

**[0039]** B) The receiver RX_IO[1] is mapped to RDATA[1] at multiplexer 161, so that both RX_IO[2] and RX_IO[1] are transmitting the same data to the multiplexer 161.

**[0040]** C) RX_IO[2] is coupled to the receive calibration circuit RXCAL 171, and RX_IO[1] is selected by multiplexer 161 to apply data to RDATA[1].

**[0041]** D) TX_IO[2] is switched to the calibration signal source TXCAL 108.

**[0042]** E) RX_IO[2] starts supplying calibration data to the receive calibration circuit RXCAL 171.

**[0043]** At the completion of these steps, the input bus TDATA[7:0] is mapped via the transmitters TX_IO[8:3,1:0] across the media 138-133, 131 and 130 to the receivers RX_IO[8:3,1:0], to output bus RXDATA[7:0]. This operation occurs without interruption in the logical layer communications. After completion of calibration on the signal path including RX_IO[2], the process waits for TX_IO[2] to begin transmitting data once again. Then, the link being calibrated is changed to the next communication line.

**[0044]** Logic on the receive and transmit sides coordinates the changing of the calibration link. One simple approach would be to provide back channel communication such as operation codes in the logical layer that coordinate synchronizing rotation of the calibration link. However, this additional complexity at the logical layer may not be necessary in some embodiments. Another approach would be to use internal counters on both sides of the link synchronized during an initialization. With sufficient timing padding around the transition points, accuracy of the synchronization requirements could be reduced allowing each side to operate essentially open loop, with the possible exception of an initialization routine which establishes a starting point.

**[0045]** Fig. 3 illustrates one process for coordinating the changing of the calibration link. On the left side of Fig. 3, a routine for the transmit side is shown. On the right side of Fig. 3, a routine for the receive side is shown. In the flowcharts, each physical communication link PHY is given the index "i". For the communication link used for calibration the index i = n.

**[0046]** Calibration starts on the transmit side at block 300. At the start, calibration data is transmitted on the physical link PHY[n], and logical data is transmitted on the physical links PHY[i] for i<n, and PHY[i+1] for i>n (block 301). This is a representative mapping of the logical data to physical links for changing the calibration link in a pattern where the calibration link changes from link 0 through link N in an increasing manner. The mapping will be adapted according to the pattern used for changing the calibration link. The transmit side waits a time interval ($\Delta T$) represented by line 302, which is long enough to allow the receive side to complete calibration. After waiting a time interval, the process switches the calibration path (block 303). According to this process, logical data for input line (n) is transmitted on PHY[n] in parallel with the transmission on PHY[n+1], for a time corresponding to a settling interval at the receiver (block 304). Then, logical data for input line (n) is transmitted only on PHY[n] (block 305). At this point, the process is ready to change the calibration link, and the index (n) is changed according to a continuous periodic pattern (block 306). Then the process loops back to block 301, and repeats.

**[0047]** On the receive side, calibration starts at block 310. To begin the process, data for calibration is received on the physical link PHY[n] (block 311). The received calibration data is processed, and the logical data signals received on the other links are routed to the receiver bus (block 312). Line maintenance or calibration is executed for PHY[n], to for example update a calibration parameter like clock phase (block 313). The receiver then waits a time interval ($\Delta T$) represented by arrow 314, to provide a margin for synchronization with the transmit side of the high-speed parallel bus. After the time interval, the receiver switches calibration path (block 315). The process to switch the calibration path includes receiving logical data for bus line (n) on both PHY[n] and PHY[n+1] (block 316). After a settling interval, PHY[n+1] is switched to the receive calibration circuit RXCAL 171, while PHY[n] is coupled to the receiver bus (block 317). Thus, the next communication link PHY[n+1] is ready to receive calibration data. The value of the index (n) is changed according to the pattern (block 318), then the process loops back to block 311, and repeats.

**[0048]** According to one embodiment of the present invention, the pattern for the continuous periodic rotation, referring to Fig. 2, would be as follows:

n = 0,1,2,3,4,5,6,7,8,7,6,5,4,3,2,1,0,1...

**[0049]** This pattern simplifies the switching between links, as the handoff occurs between adjacent links in each step. However, the time between updates is not the same for all links. This difference in time between updates may not be a problem for a small number of links. However, larger error terms may be encountered for the communication systems in which many blocks are daisy chained together.

**[0050]** The worst-case update rate is described in equation 1, where N is the number of actual data links, and $T_{cal}$ is the total time for a link to calibrate and handoff to the next link.

Eq. 1) $$T_{update} = 2*N*T_{cal}$$

**[0051]** Assuming a link takes about 1000 cycles to calibrate a clock at 400 MHz, and adding some time for synchronization, $T_{cal}$ can be approximated as five microseconds. This means the update frequency for an eight link system would be around 12.5 kHz. A reduction of calibration time by factor of 10 could potentially increase update frequency to about 150 kHz, if such a scheme could maintain synchronization and have the accuracy necessary for a given application. Other items such as spread spectrum clocking may affect the desired update rate.

**[0052]** By way of example, Fig. 4 is a simplified diagram of a system suitable for use with the continuous periodic calibration technique of the present invention, where the calibration is applied to adjusting clock phase to select optimal sampling point for a physical channel. In this example, only the receive side is illustrated. However, both transmit and receive sides receive a reference clock, in this example a 400MHz clock on line 400. The reference clock is applied to a phase-locked loop circuit 401 which multiplies a clock by eight in this example to produce a receive clock at a frequency of 3.2GHz. One copy of the receive clock is applied to each of the receivers, such as receiver 141. The receiver 141 includes a clock phase adjustment circuit 402 which

applies a clock at the calibrated sampling point to receive sense amplifier 403. Input from the physical channel at 6.4 gigabits per second, where sampling occurs on each transition of clock is received through a buffer 404 into the sense amplifier 403. The output of the sense amplifier 403 is applied to the calibration circuit, which comprises the mixer 405, and a source 406 of a pseudorandom bit sequence used for calibration. The received pseudorandom bit sequence RX_PRBS(n) on the physical layer is applied to the mixer 405. The mixer 405 produces an adjustment parameter on line 407, which is applied to the clock phase adjuster 402. When the output of the sense amplifier 403 is used for data, it is applied to a serial-to-parallel converter 408, clocked at the reference clock rate, e.g. 400MHz, to apply a parallel output at the reference clock rate. In this example, translating 6.4 gigabits per second to a 400MHz clock would involve an 8-bit serial-to-parallel converter 408 for each of the receive signal paths.

**[0053]** Although much of the discussion has been addressed to calibration from the perspective of timing calibration, the invention can be applied to other functions requiring periodic adjustment of the physical channel for data transmission, such as current calibration, resistor calibration, adaptive equalization, and other types of line maintenance and tuning.

**[0054]** In some embodiments, a hybrid method of sending a source synchronous clock, along with continuous periodic calibration as described above can be used at the expense of an additional physical layer link to carry the clock. In Fig. 5, the system of Fig. 1 is shown using the same reference numerals, with an additional physical layer link 500 from a source clock 501 on the first integrated circuit 10 to a clock circuit 502 on the second integrated circuit used for carrying a source synchronous clock, which may be desirable in some environments where source synchronous clocking provides superior performance, and can be used in combination with the continuous periodic calibration process described herein. The additional physical layer link 500 carrying the clock may or may not be included in the continuous periodic calibration routine, as suits a particular implementation.

**[0055]** According to the present invention, continuous periodic calibration provides a solution for tracking slowly changing drift terms for ideal sampling of clock relative to

data. It seamlessly calculates the best sampling point of data continuously, with the overhead of one extra IO, without knowledge of the logical layer.

**[0056]** The technique is particular suited to high-speed chip to chip communications.

**[0057]** In summary, the present invention provides methods and apparatus for providing continuous calibration of properties associated with a parallel interface that includes N links. The calibrated property may include for example signaling levels, optimal placement of sampling times for symbol capture, and impedance of the termination element or equalization coefficients associated with one of the N links. In one embodiment, continuous calibration of an optimal timing point for sampling by receiver circuit is provided using an additional link (N+1) to time multiplex a calibration sequence among the N links. The calibration is rotated or switched among the N+1 links, while normal communication is executed on the other N links.

**[0058]** While the present invention is disclosed by reference to the preferred embodiments and examples detailed above, it is to be understood that these examples are intended in an illustrative rather than in a limiting sense. It is contemplated that modifications and combinations will readily occur to those skilled in the art, which modifications and combinations will be within the spirit of the invention and the scope of the following claims. What is claimed is: